



Cite this: *Phys. Chem. Chem. Phys.*,
2021, **23**, 4736

Efficient and low-scaling linear-response time-dependent density functional theory implementation for core-level spectroscopy of large and periodic systems†

Augustin Bussy * and Jürg Hutter*

We discuss our implementation of linear-response time-dependent density functional theory (LR-TDDFT) for core level near-edge absorption spectroscopy. The method is based on established LR-TDDFT approaches to X-ray absorption spectroscopy (XAS) with additional accurate approximations for increased efficiency. We validate our implementation by reproducing benchmark results at the K-edge and showing that spin-orbit coupling effects at the $L_{2,3}$ -edge are well described. We also demonstrate that the method is suitable for extended systems in periodic boundary conditions and measure a favorable sub-cubic scaling of the calculation cost with system size. We finally show that GPUs can be efficiently exploited and report speedups of up to a factor 2.

Received 27th November 2020,
Accepted 7th February 2021

DOI: 10.1039/d0cp06164f

rsc.li/pccp

X-ray absorption spectroscopy (XAS) is a very useful tool to investigate the structure of matter. It is a local and element specific probe, providing insights into both geometric and electronic properties of molecules and solids. Advances in experimental techniques and increased availability of high quality X-ray sources have turned XAS into a central characterization technique in many fields, from surface sciences to biochemistry. The high energy of X-rays allow excitations of electrons from core states, yielding information otherwise inaccessible.

The XAS spectrum can be split into two regions, first of which is the X-ray absorption near edge structure (XANES), situated close to the ionization edge and mostly yielding insights about the electronic structure. At higher energies, the extended X-ray absorption fine structure (EXAFS) is found, providing more geometrical information. Since XAS is element specific, multiple spectra can be resolved for a given sample, one for each atomic kind and core state combination. The nomenclature depends on the donor core state, namely K-edge for 1s excitations, L_1 -edge for 2s, $L_{2,3}$ -edge for 2p, *etc.*

In order to help interpret XANES experiments, significant theoretical efforts were made leading to the development and adaptation of a variety of simulation methods. Well known

density functional theory (DFT) based methods such as Δ SCF¹ and transition-potential DFT^{2,3} deliver insights at an affordable computational cost. The former is however only well defined for the first excited state, making the simulation of a full spectrum challenging (but possible⁴). Transition-potential DFT methods do not suffer from this, although the spectra they produce are typically of lower quality than the next level of theory: time-dependent DFT. There are two approaches to it, real-time TDDFT⁵ and linear-response TDDFT^{6–8} (LR-TDDFT), both commonly used for core electronic spectroscopy. However, the improved results come at the expense of more complex implementations and increased computational costs. Much more involved are wave function based approaches such as those from the coupled-cluster family.^{9,10} Spectra produced by such methods are usually in excellent agreement with experiments, but their high computational cost make them prohibitively expensive for everything but small molecules.¹¹

We believe that LR-TDDFT with hybrid functionals offers a good compromise between cost and accuracy for XANES simulations. Throughout this work, we report our implementation of it in the CP2K¹² software package. Building on common approaches to XAS LR-TDDFT, we include additional core-specific cost-reduction measures and prove their validity. We also show that our method can accurately describe the $L_{2,3}$ -edge where spin-orbit effects are important and demonstrate its applicability in periodic boundary conditions. We finally discuss the favorable scaling of our method and how GPUs can be exploited for additional speed.

Department of Chemistry, University of Zurich, Winterthurerstrasse 190, CH-8057 Zurich, Switzerland. E-mail: augustin.bussy@chem.uzh.ch, hutter@chem.uzh.ch

† Electronic supplementary information (ESI) available. See DOI: 10.1039/d0cp06164f



1 Theory

1.1 Linear-response TDDFT

For a system of N interacting electrons in the ground state, the time-independent Kohn–Sham (KS) equations

$$F^\sigma |\phi_{i\sigma}^0\rangle = \varepsilon_{i\sigma} |\phi_{i\sigma}^0\rangle \quad (1)$$

have to be solved for the KS orbitals $\phi_{i\sigma}^0(\mathbf{r})$ and corresponding orbital energies $\varepsilon_{i\sigma}$, where σ labels the spin and $N = N_\alpha + N_\beta$. The KS Hamiltonian is expressed as

$$F^\sigma(\mathbf{r}) = -\frac{1}{2}\nabla^2 + V_{\text{eff}}^\sigma(\mathbf{r}) \quad (2)$$

where the effective single-particle potential is the sum of an external, Hartree and exchange–correlation potentials:

$$V_{\text{eff}}^\sigma(\mathbf{r}) = V_{\text{ext}}^\sigma(\mathbf{r}) + \int d\mathbf{r}' \frac{n(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} + \frac{\delta E_{\text{xc}}[n_\alpha, n_\beta]}{\delta n_\sigma} \quad (3)$$

where $n(\mathbf{r}) = n_\alpha(\mathbf{r}) + n_\beta(\mathbf{r})$ is the electronic density and $n_\sigma(\mathbf{r}) = \sum_i |\phi_{i\sigma}^0(\mathbf{r})|^2$. The exchange–correlation potential above is defined as the functional derivative of the approximate exchange–correlation energy functional, expressed in hybrid DFT as

$$E_{\text{xc}}[n] = E_{\text{c}}^{\text{DFT}}[n] + (1 - c_{\text{HF}})E_{\text{x}}^{\text{DFT}}[n] + c_{\text{HF}}E_{\text{x}}^{\text{HF}}[n] \quad (4)$$

with a fraction c_{HF} of Hartree–Fock exchange.

In practice, the ground state KS orbitals are expanded in a finite basis set $\{\varphi_p(\mathbf{r})\}$, which is here taken to be made of non-orthogonal atom-centered Gaussian type orbitals (GTOs), such that

$$\phi_{i\sigma}^0(\mathbf{r}) = \sum_p c_{pi\sigma}^0 \varphi_p(\mathbf{r}) \quad (5)$$

and

$$S_{pq} = \langle \varphi_p | \varphi_q \rangle. \quad (6)$$

The KS equations can then be cast into matrix form as an Hermitian generalized eigenvalue problem, to be solved for the coefficients $\{c_{pi\sigma}^0\}$:

$$\mathbf{F}^\sigma \mathbf{c}^0 = \mathbf{S} \mathbf{c}^0 \varepsilon \quad (7)$$

where ε is a matrix with the orbital energies on the diagonal and $F_{pq}^\sigma = \langle \varphi_p | F^\sigma | \varphi_q \rangle$. The resulting ground state KS orbitals and energies, obtained through a self-consistent field (SCF) procedure, serve as a base for LR-TDDFT. Note that throughout this work, KS orbitals and molecular orbitals (MOs), as well as atom-centered Gaussian basis functions and atomic orbitals (AOs), will be referred to interchangeably.

For the simulation of electronic absorption spectroscopy in the Sternheimer approach to LR-TDDFT,^{13,14} the absorbed photon is modelled as a harmonic perturbation to the external ground state potential:

$$V_{\text{ext}}(t, \mathbf{r}) = V_{\text{ext}}^0(\mathbf{r}) + V^-(\mathbf{r})e^{-i\omega t} + V^+(\mathbf{r})e^{i\omega t} \quad (8)$$

which, in virtue of the fundamental theorems of TDDFT,^{15,16} leads to a similar perturbation to the ground state density and

KS orbitals:

$$n_\sigma(t, \mathbf{r}) = n_\sigma^0(\mathbf{r}) + n_\sigma^-(\mathbf{r})e^{-i\omega t} + n_\sigma^+(\mathbf{r})e^{i\omega t} \quad (9)$$

$$\phi_{i\sigma}(t, \mathbf{r}) = \phi_{i\sigma}^0(t, \mathbf{r}) + \phi_{i\sigma}^-(\mathbf{r})e^{-i\omega t} + \phi_{i\sigma}^+(\mathbf{r})e^{i\omega t}$$

where $\phi_{i\sigma}^\pm(\mathbf{r})$ are the so-called response orbitals, which are orthogonal to the ground state KS orbitals $\langle \phi_{i\sigma}^0 | \phi_{j\tau}^\pm \rangle = 0$ and assumed to be real. These perturbed quantities can be plugged in the time-dependent KS equations

$$\left[-\frac{1}{2}\nabla^2 + V_{\text{eff}}^\sigma(t, \mathbf{r}) \right] \phi_{i\sigma}(t, \mathbf{r}) = i \frac{\partial}{\partial t} \phi_{i\sigma}(t, \mathbf{r}) \quad (10)$$

and only the first order terms kept. After expanding the ground state KS orbitals and the linear response orbitals into the atom-centered non-orthogonal basis $\{\varphi_p(\mathbf{r})\}$ as in eqn (5) and (6) and

$$\phi_{i\sigma}^\pm(\mathbf{r}) = \sum_p c_{pi\sigma}^\pm \varphi_p(\mathbf{r}), \quad (11)$$

the vertical excitation energies of the system are obtained by solving the non-Hermitian eigenvalue problem:

$$\omega \begin{pmatrix} -\mathbf{G} & \mathbf{0} \\ \mathbf{0} & \mathbf{G} \end{pmatrix} \begin{pmatrix} \mathbf{c}^+ \\ \mathbf{c}^- \end{pmatrix} = \begin{pmatrix} \mathbf{A} + \mathbf{B} - \mathbf{D} & \mathbf{B} - \mathbf{E} \\ \mathbf{B} - \mathbf{E} & \mathbf{A} + \mathbf{B} - \mathbf{D} \end{pmatrix} \begin{pmatrix} \mathbf{c}^+ \\ \mathbf{c}^- \end{pmatrix} \quad (12)$$

where the metric \mathbf{G} is defined as

$$G_{pi\sigma, qj\tau} = S_{pq} \delta_{ij} \delta_{\sigma\tau}, \quad (13)$$

the ground state information is contained in matrix \mathbf{A}

$$A_{pi\sigma, qj\tau} = (F_{pq}^\sigma - \varepsilon_i S_{pq}) \delta_{ij} \delta_{\sigma\tau}, \quad (14)$$

\mathbf{B} is the Hartree exchange–correlation kernel matrix (integrals in Mulliken notation)

$$B_{pi\sigma, qj\tau} = \sum_{rs} Q_{pr}^\sigma \left(r i_\sigma \left| f_{\sigma,\tau}^{Hxc} \right| s j_\tau \right) Q_{sq}^\tau, \quad (15)$$

and \mathbf{D} , \mathbf{E} are the on- and off-diagonal Hartree–Fock exchange kernel matrices, respectively

$$D_{pi\sigma, qj\tau} = c_{\text{HF}} \delta_{\sigma\tau} \sum_{rs} Q_{pr}^\sigma (r s | i_\sigma j_\tau) Q_{sq}^\tau \quad (16)$$

$$E_{pi\sigma, qj\tau} = c_{\text{HF}} \delta_{\sigma\tau} \sum_{rs} Q_{pr}^\sigma (r j_\tau | s i_\sigma) Q_{sq}^\tau \quad (17)$$

The matrix \mathbf{Q}^σ introduced above is a projector on the subspace of the unoccupied unperturbed states and is defined as

$$Q_{pq}^\sigma = \left\langle \varphi_p \left| \left(1 - \sum_i \phi_{i\sigma}^0 \phi_{i\sigma}^0 \right) \right| \varphi_q \right\rangle = (1 - \mathbf{S} \mathbf{P}^\sigma)_{pq} \quad (18)$$

with the density matrix $P_{pq}^\sigma = \sum_i c_{pi\sigma}^0 c_{qi\sigma}^0$. The Hartree exchange–correlation kernel in eqn (15) is, within the adiabatic approximation,¹⁷ expressed as

$$f_{\sigma,\tau}^{Hxc}(\mathbf{r}, \mathbf{r}') = \frac{1}{|\mathbf{r} - \mathbf{r}'|} + \frac{\delta^2 E_{\text{xc}}}{\delta n_\sigma(\mathbf{r}) \delta n_\tau(\mathbf{r}')} \Big|_{n^0}. \quad (19)$$

Note that the exchange part of the exchange–correlation kernel is scaled by $(1 - c_{\text{HF}})$ for consistency. All four-center



two-electron integrals needed involve both MOs and AOs.

Non-Hermitian eigenvalue problems such as eqn (12) are harder and more expensive to solve than their Hermitian counterparts. Casida¹⁸ showed that, assuming real orbitals, the LR-TDDFT eigenvalue problem can be cast into Hermitian form at the cost of a matrix square root. Another way of simplifying the problem is the so-called Tamm–Dancoff approximation (TDA),¹⁹ which yields excitation energies of similar quality with respect to full TDDFT.²⁰ The TDA is obtained by setting $\mathbf{c}^+ = 0$ in eqn (12), resulting in

$$\omega \mathbf{G} \mathbf{c}^- = (\mathbf{A} + \mathbf{B} - \mathbf{D}) \mathbf{c}^-, \quad (20)$$

which has the multiple advantages of turning the eigenvalue problem Hermitian, dropping the off-diagonal exchange kernel matrix \mathbf{E} and reducing the size of the matrices by a factor 2. Note that in the case of restricted closed-shell systems, where $\phi_{i\alpha}^0(\mathbf{r}) = \phi_{i\beta}^0(\mathbf{r})$, the problem can be simplified further by treating singlet and triplet excitations independently.²¹ The spin index σ is then dropped and the matrix dimensions halved again. Since triplet excitations from a singlet ground state are spin-forbidden in the dipole approximation, it is usually sufficient to solve the singlet problem only.

1.2 Core specific approximations

Obtaining core excitation energies from standard LR-TDDFT is impractical at best since the corresponding eigenvalues, which are much larger than those of valence excitations, are not efficiently computed by conventional iterative solvers. Therefore, extra steps need to be taken to model excitation of core electrons in an affordable manner.

In 1980, in the context of many-body theory of core holes, Cederbaum *et al.*²² argued that core and valence states are only weakly coupled due to their large energy difference and localization in space. This was later adapted to LR-TDDFT by Stener *et al.*,⁶ which allows for ignoring valence excitations when building the eigenvalue problem. In practice, this reduces the span of the MO indices i, j in matrices (13–17) to only a handful of core states, drastically reducing their size. This also brings core to low lying unoccupied bound states excitation to the smallest eigenvalues, making iterative solvers relevant again. This approach is commonly used in the community,^{5,7,23} although other methods based on core specific iterative solvers are also available.^{24,25}

DeBeer George *et al.*⁸ went one step further by invoking the sudden approximation,²⁶ which neglects electronic relaxation of the system beyond the core region after a core electron is excited. Combined with the localized nature of core states, this allows further decoupling of core excitations. The XANES spectrum can then be built as a sum of independent one donor core orbital contributions. This approach boosts the efficiency even further as diagonalizing a series of small matrices scales better than solving one large eigenvalue problem. This however requires donor core MOs to be local in space, which is not necessarily the case for canonical KS orbitals that are equivalent under symmetry. In such instances, the core orbitals need to be

actively localized using schemes such as Foster and Boys,²⁷ Pipek and Mezey²⁸ or maximally localized Wannier functions.²⁹ Note that in the case of $L_{2,3}$ spectroscopy, the three degenerate 2p states must be treated simultaneously.

With the matrix sizes and diagonalization costs brought down, the bottleneck of such TDDFT calculations may become the evaluation of 4-center 2-electron integrals needed for the kernel matrices (15), (16) and (17). The cost of these integrals formally scales with the fourth power of the system size, although symmetries and resolution of the identity (RI) schemes³⁰ are usually used for efficiency. In the context of XANES, all electron repulsion integrals involve a core orbital, which locality can be exploited for additional screening.³¹ We further utilize the local nature of core states by introducing a reduced RI scheme for increased efficiency. For the Hartree kernel:

$$(p_i \sigma | q_j \tau) \approx \sum_{\mu\nu} (p_i \sigma | \mu) (\mu | \nu)^{-1} (\nu | q_j \tau) \quad (21)$$

where the atom-centered basis functions $\varphi_p(\mathbf{r})$, $\varphi_q(\mathbf{r})$ and the core orbitals $\phi_{i\sigma}(\mathbf{r})$, $\phi_{j\tau}(\mathbf{r})$ must overlap for non-zero products. This allows for a reduced RI basis $\{\chi_{\mu}(\mathbf{r})\}$ only made of Gaussian functions centered on the excited atom. Therefore, this integral is a purely local quantity that is very efficient to evaluate. The exchange–correlation kernel integrals are treated in a similar way:

$$\begin{aligned} (p_i \sigma | f_{\sigma,\tau}^{\text{xc}} | q_j \tau) &\approx \sum_{\kappa\lambda\mu\nu} (p_i \sigma | \kappa) (\kappa | \lambda)^{-1} (\lambda | f_{\sigma,\tau}^{\text{xc}} | \mu) \\ &\times (\mu | \nu)^{-1} (\nu | q_j \tau) \end{aligned} \quad (22)$$

where the 3-center $(p_i \sigma | \kappa)$ and inverse 2-center $(\kappa | \lambda)^{-1}$ integrals are the same as in eqn (21). Because the RI basis functions $\chi_{\mu}(\mathbf{r})$, $\chi_{\nu}(\mathbf{r})$ are centered on the excited atom, it is sufficient to evaluate $f_{\sigma,\tau}^{\text{xc}}$ in its vicinity only. To that end, the electronic density is projected on the RI basis

$$\begin{aligned} n(\mathbf{r}) &= \sum_{\sigma} \sum_{pq} P_{pq}^{\sigma} \varphi_p(\mathbf{r}) \varphi_q(\mathbf{r}) \\ &\approx \sum_{\sigma} \sum_{pq} \sum_{\mu\nu} P_{pq}^{\sigma} (pq\mu) S_{\mu\nu}^{-1} \chi_{\nu}(\mathbf{r}) \\ &= \sum_{\nu} d_{\nu} \chi_{\nu}(\mathbf{r}), \end{aligned} \quad (23)$$

turning it into a linear combination of RI basis functions and making the numerical integration of $(\lambda | f_{\sigma,\tau}^{\text{xc}} | \mu)$ straight forward. Note that the contributions from core electrons of neighboring atoms might not be caught well, as Gaussian functions lack sharpness away from their origin. This is addressed by either adding extra basis elements centered on first neighbors for the projection, or using pseudopotentials for their description. The Hartree–Fock exchange kernel requires integrals of the type

$$(p_i \sigma | i_{\sigma} j_{\tau}) \approx \sum_{\mu\nu} (pq\mu) (\mu | \nu)^{-1} (\nu | i_{\sigma} j_{\tau}) \quad (24)$$

which benefit from the same RI scheme. They are, however, much more computationally demanding since all overlapping atomic centered basis functions $\varphi_p(\mathbf{r})$, $\varphi_q(\mathbf{r})$ in the system contribute. For large molecules and periodic systems, the cost can be limited by



using short range exchange operators such as the smooth error function³² or the sharper truncated Coulomb potential.³³

After solving eqn (12), the XANES spectrum is obtained by calculating oscillator strengths in the dipole approximation. LR-TDDFT is known to produce correct spectra, with accurate feature spacing and relative intensities, but in the wrong place on the energy axis. This is mainly due to the lack of relaxation upon the core hole creation and self interaction error. Different solutions have been proposed to tackle the latter,^{34–36} including new core excitation specific range-separated hybrid functionals.^{37,38} However, when standard XAS LR-TDDFT is employed, the calculated spectra usually have to be shifted to match experiments, either empirically or using a Δ SCF calculation.^{11,39,40}

We implemented our method in the open source CP2K¹² software package, based on the GAPW method⁴¹ for ground state calculations. By default, the RI basis sets are automatically generated using the method of Stoychev and co-workers.⁴² The 2- and 3-center Coulomb and exchange integrals are computed analytically using the Libint library⁴³ and the exchange–correlation kernel is evaluated using the LIBXC⁴⁴ library. Most basis sets come from Basis Set Exchange.⁴⁵

1.3 Spin–orbit coupling

Relativistic effects play a role when dealing with core electrons as they attain very high speeds, typically reaching considerable fractions of the speed of light. The simplest and most well known way of treating such effects is the scalar relativistic approximation, in which only the mass-velocity and Darwin terms of Dirac's equations are kept. This however leads to the complete neglect of the spin–orbit coupling (SOC), which is at the origin of many important effects in chemistry and physics, such as L_{2,3}-edge peak splitting in XANES.

Within LR-TDDFT, scalar relativistic effects are added at the ground state level as a modification to the KS Hamiltonian, lowering the orbital energies of core states. Thus, resulting spectra are blue-shifted with respect to non-relativistic calculations, but the overall shape remain essentially unchanged. Since XAS LR-TDDFT spectra need to be rigidly shifted anyways, scalar relativistic effects can usually be assumed to be accounted for that way.¹¹ Adding SOC to TDDFT is less trivial and generally involves more computationally intensive methods such as two-^{56,57} or even four-component⁵⁸ relativistic TDDFT. A simpler approach consists of adding SOC as a perturbation to a non-relativistic calculation by coupling excited state wave functions of different spin multiplicities.^{59,60} This can be rigorously done within LR-TDDFT using the auxiliary set of many-electron wave functions (AMEW) framework proposed by de Carvalho and co-workers.⁶¹

In the general case of the LR-TDDFT Sternheimer based formalism, the AMEW expansion for an excited state I with excitation energy ω_I is:

$$|\Psi_I\rangle = \sum_{\{+,-\}} \sum_i^{N_o} \sum_{\sigma} (r_{i\sigma}^{I\pm})^\dagger \hat{a}_{i\sigma} |\Psi_0\rangle \quad (25)$$

where $|\Psi_0\rangle$ is the Slater determinant (SD) of all occupied KS orbitals, $\hat{a}_{i\sigma}$ is the KS orbital annihilation operator and $(r_{i\sigma}^{I\pm})^\dagger$ is

the creation operator for the LR orbitals. Note that this expression simplifies for core excitation within the sudden approximation as the sum over occupied states reduces to one core state only (three for degenerate 2p states). In case of TDA, the sum over $\{+, -\}$ vanishes as the $c_{pi\sigma}^+$ LR coefficients are set to zero.

Following quasi-degenerate perturbation theory (QDPT),⁵⁹ the Born–Oppenheimer Hamiltonian is perturbed with the addition of SOC matrix elements

$$H_{ISM,JS'M'} = \delta_{IJ}\delta_{SS'}\delta_{MM'}\omega_I^{SM} + \langle \Psi_I^{SM} | H_{SO} | \Psi_J^{S'M'} \rangle \quad (26)$$

where $|\Psi_I^{SM}\rangle$ is the AMEW expansion of an excited state with energy ω_I^{SM} and spin quantum numbers $s = S$ and $m_s = M$. H_{SO} is the spin–orbit Hamiltonian, chosen here to be the ZORA Hamiltonian⁶² together with van Wüllen's model potential.⁶³ Note that because of its one-electron character, SOC matrix elements of any two states' SDs can simply be evaluated using Löwdin's rule.⁶⁴ The detailed matrix elements are given in the ESI.†

In restricted closed-shell systems, four kinds of excited states can cohabit; singlets with quantum numbers $s = 0$, $m_s = 0$, spin-conserving triplets with $s = 1$, $m_s = 0$ and spin–flip triplets with $s = 1$, $m_s = \pm 1$. All of them can be described within the AMEW formalism and standard LR-TDDFT. Diagonalizing the complex Hermitian matrix yields new excited states as linear combination of singlets and triplets AMEWs with excitation energies that are corrected for SOC. The above formalism can easily be extended to open-shell systems by adding AMEWs built on spin–flip TDDFT^{65–67} to the standard spin-conserving excitations. Note that systems with a strong multi-reference character would still be out of reach, as LR-TDDFT is inherently a single reference method.

Adding SOC effects for L-edge spectroscopy does not come without cost. Indeed, the TDDFT eigenvalue problem needs to be solved twice (once for singlets and once for triplets) and the QDPT matrix can potentially reach very large sizes, leading to expensive diagonalizations. That latter point can easily be addressed by only considering excited states lying in the near-edge region of the spectra, which are limited in number.

1.4 ADMM acceleration

With the efficient core specific approximations described above, the bottleneck of XAS LR-TDDFT calculation may actually become the ground state KS SCF procedure, especially at the hybrid DFT level. Indeed, the addition of non-local Hartree–Fock exchange (HFX) to local generalized gradient approximation (GGA) adds a heavy computational burden that technically scales with the fourth power of the system size, as 4-electron 2-center electron repulsion integrals (ERIs) need to be evaluated. Although integral screening⁶⁸ and the use of short range exchange operators^{32,69} have done much to bring that cost down, scaling with respect to basis set quality remains poor. To tackle this issue, the auxiliary density matrix method (ADMM) was developed⁷⁰ and implemented in the CP2K software package. The HFX energy is a functional of the density matrix P



and can be written as

$$E_X^{\text{HF}} = E_X^{\text{HF}}[\hat{P}] + (E_X^{\text{HF}}[P] - E_X^{\text{HF}}[\hat{P}]) \quad (27)$$

$$\approx E_X^{\text{HF}}[\hat{P}] + (E_X^{\text{GGA}}[P] - E_X^{\text{GGA}}[\hat{P}])$$

where the auxiliary density matrix $\hat{P} \approx P$ was introduced. The assumption is that the difference in Hartree–Fock exchange energy between the primary and the auxiliary density matrices is well captured by GGA, which allows for the costly evaluation of ERIs in a smaller and/or less diffuse auxiliary basis set.

There are two main flavors of ADMM, known as purified wave function fitting or ADMM1 and nonpurified wave function fitting or ADMM2. In the latter approach, the auxiliary MO coefficients are obtained by minimizing

$$\sum_i \int (\psi_i(\mathbf{r}) - \hat{\psi}_i(\mathbf{r}))^2 d\mathbf{r}, \quad (28)$$

which is the square difference between the occupied wave functions in the primary and auxiliary basis sets, respectively. The auxiliary density matrix comes from the simplest projection of the MO coefficients from one basis to the other. On the other hand, purified wave function fitting minimizes (28) with the additional constraint that $\psi_i(\mathbf{r})$ and $\hat{\psi}_i(\mathbf{r})$ be orthogonal. This ensures that the auxiliary density matrix fulfills the properties of a pure density matrix, *i.e.* symmetry, idempotency and particle conservation. Note that, in purified wave function fitting, the ADMM Kohn–Sham matrix resulting from the SCF procedure cannot directly serve as a base for LR-TDDFT, as its eigenvalue do not correspond to orbital energies. As described in the original ADMM paper, an additional step for eigenvalue correction needs to be taken. Nonpurified wave function fitting, ADMM2, does not require any post SCF treatment. The exact speedup in HFX energy evaluation due to the ADMM method depends on the choice of primary and auxiliary basis sets and can typically reach orders of magnitude. Further studies have

confirmed that the assumption expressed in eqn (27) is sound⁷¹ and that ADMM is usually faster, although less accurate, than other RI based acceleration schemes.⁷²

2 Validation

2.1 K-edge

To assess the quality of the method, the first excitation energy at the K-edge was computed for a series of small molecules for a variety of functionals and basis sets (without TDA). The structures were optimized at the B3LYP/cc-pVTZ level. Our results, displayed in Table 1, are in very good agreements with other similar publications (ref. 78 and 79). In particular, the mean absolute deviation (MAD) of the energies obtained with the cc-pCVTZ basis set is only 0.02 eV compared to those of Inanura, Otsuka and Nakai.⁷⁸ This is significant since they solved the full LR-TDDFT equations, without any further approximation. The MAD compared to the work of Lestrangé, Nguyen and Li,⁷⁹ who use a energy specific iterative eigensolver for core excitations and the 6-311+G** basis set, is 0.05 eV. The small amplitude of the MADs indicates that the introduction of the sudden approximation and our local RI scheme have a negligible impact on accuracy. Note that the calculations of these averages do not involve all entries of Table 1, since the overlap with the other works is not complete. Experimental excitation energies are not reported here as the focus is on the consistency of our method with respect to other TDDFT implementations only.

Additionally, basis set convergence studies were conducted at the K-edge. The difference between the first and second excitation energies ($\omega_2 - \omega_1$) was computed for the H₂O, CO, NH₃ and C₂H₄ molecules across the members of the pcseg⁸⁰ and the core-specific pcX⁸¹ basis families. Note that the converging quantity was chosen such that self-interaction errors and related shifts do not play a role. The energy difference was

Table 1 K-edge first excitation energies for small molecules, using the all-electron Gaussian basis sets cc-pCVTZ^{46,47} and 6-311+G**.^{48,49} The bold and italic atom in the chemical formula indicates the excited atom. For nitrous oxide, N^T represents the terminal nitrogen and N^C the central nitrogen. The standard GGA functionals BLYP^{50,51} and PBE⁵² as well as hybrid functionals with increasing fraction of Hartree–Fock exchange B3LYP (20%),⁵³ PBE0 (25%)⁵⁴ and BHHLYP (50%)⁵⁵ were used

	cc-pCVTZ					6-311+G**				
	BLYP	PBE	B3LYP	PBE0	BHHLYP	BLYP	PBE	B3LYP	PBE0	BHHLYP
C ₂ H ₂	269.97	269.26	275.36	276.29	283.68	269.74	268.95	275.83	276.16	283.62
C ₂ H ₄	269.05	268.29	274.36	275.24	282.58	268.82	268.07	274.24	275.13	282.53
CO	271.32	270.56	276.17	276.90	283.58	271.27	270.41	276.10	276.84	283.57
CH ₂ O	270.23	269.40	275.26	275.98	283.01	270.09	269.24	275.17	275.90	282.96
CHFO	272.69	271.83	277.67	278.35	285.28	272.57	271.70	277.61	278.29	285.25
CF ₂ O	275.43	274.53	280.31	280.93	287.76	275.29	274.40	280.19	280.84	287.71
N ₂	382.49	381.77	388.55	389.70	397.97	382.37	381.61	389.32	389.64	397.98
NH ₃	380.97	380.47	388.09	389.74	399.08	380.67	380.04	388.02	389.45	398.96
N ^T NO	382.59	381.89	388.89	390.14	398.68	382.40	381.68	388.81	390.05	398.66
NN ^C O	386.16	385.46	392.33	393.53	401.86	385.96	385.24	392.16	393.36	401.73
CO	512.33	511.63	519.82	521.45	531.65	512.12	511.43	519.71	521.35	531.63
N ₂ O	512.49	511.83	520.46	522.25	533.14	512.29	511.61	520.34	522.14	533.09
CH ₂ O	509.41	508.70	516.68	518.23	528.13	509.24	508.53	516.60	518.16	528.13
CHFO	510.24	509.53	517.64	519.23	529.28	510.08	509.37	517.54	519.15	529.24
CF ₂ O	510.96	510.26	518.44	520.07	530.21	510.79	510.06	518.31	519.94	530.12
CHFO	659.95	659.17	669.59	671.76	684.99	659.79	659.00	669.47	671.65	684.92
CF ₂ O	661.36	660.59	671.02	673.21	686.43	661.12	660.35	670.83	673.04	686.29



generally found to converge quickly. For H₂O, NH₃ and C₂H₄, triple zeta quality basis sets were already sufficient. It was also observed that using the core uncontracted pcX basis sets did not drastically change the convergence rate nor the converged values. The detailed results can be found in the ESI.†

2.2 L-edge

At the L-edge, spin-orbit coupling is responsible for the splitting of the L_{2,3} peaks. To demonstrate the validity of our method together with the AMEW formalism and QDPT, we simulated the L_{2,3}-edge XANES spectra for three closed-shell molecules and transition metal complexes, namely TiCl₄, SiCl₄ and CrO₂Cl₂ (without TDA). The results are displayed together with experimental data on Fig. 1. The computed spin-orbit splitting is remarkably consistent with experimental results and the spectral features are generally well reproduced, both in relative intensities and relative energies. Note that the choice of basis set and functional is such that our results can be directly compared to those of Kasper and co-workers,⁵⁶ who use a more advanced two-component relativistic TDDFT scheme. Despite the comparative simplicity of our method, the spectra we obtain are very similar. To illustrate the importance of SOC in L_{2,3}-edge calculations, the second row in Fig. 1 shows LR-TDDFT results without the AMEW and QDPT treatment. Except for arguably SiCl₄ where spin-orbit splitting is small, leading to overlapping L₂ and L₃ features, the simulated spectra have neither descriptive nor predictive power.

2.3 ADMM acceleration

The performance of the ADMM1 and ADMM2 schemes have been assessed using the CORE65 benchmark set⁸² structures, which contains 32 small first and second row molecules. The set is initially meant for core ionization benchmarks, covering 30 C1s, 21 O1s, 11 N1s and 3 F1s core levels, although we use it here for core excitations. Reference full HFX calculations were

done over the whole set using hybrid functionals with increasing fractions of exact exchange, namely B3LYP (20%), PBE0 (25%), PBEh⁸³ with $\alpha = 0.45$ (45%) and BHHLYP (50%). The basis sets of choice were the pcseg (double-, triple- and quadruple-zeta quality) and corresponding admm⁸⁴ families, which are, to the best of our knowledge, the only available all-electron ADMM compatible basis sets. Results, computed within the TDA, are displayed on Fig. 2.

The mean absolute deviation of the first excitation energies, shown as clear wide bars, varies significantly depending on the basis set and functional combination, from 0.03 eV for pcseg-3/PBE0 up to 0.23 eV for pcseg-2/BHHLYP. Logically, the error systematically increases with the Hartree-Fock fraction of the functional. Basis set quality, however, does not seem to correlate with the error.

Arguably more important than the absolute first excitation energy is the energy spacing between spectral features, since the LR-TDDFT spectrum needs shifting anyways. The thin dark bars represent the MAD of the first and third excited energies difference, *i.e.* $\omega_3 - \omega_1$, over the different basis sets and functionals. The error systematically decreases with basis set quality and, from triple-zeta on, is better than the absolute ω_1 error. Note that the somewhat arbitrary $\omega_3 - \omega_1$ difference was chosen as to avoid ω_2 , which is often degenerate with the first excitation energy.

There is neither significant nor systematic difference between the two ADMM schemes. Since the nonpurified flavor (ADMM2) is simpler and requires less operations, this makes it the go to choice, especially for larger calculations. Note that finally, all errors reported here are quite small compared to the span of the near-edge region of the XAS spectrum which typically stretches over 15 to 20 eV. Therefore, inaccuracies due to the introduction of ADMM for the ground state should be barely noticeable. The speedups measured for HFX 4-center integral evaluations were 10.1, 8.6 and 3.8 for pcseg-1, pcseg-2 and pcseg-3, respectively.




Fig. 1 Top row: Simulated XANES L_{2,3}-edge spectra for (a) TiCl₄, (b) SiCl₄ and (c) CrO₂Cl₂. Calculations are done at the B3LYP/def2-TZVPD^{73,74} level with SOC added on top of XAS LR-TDDFT, combining the AMEW and QDPT formalisms. The calculated spectra were rigidly shifted by 10.44, 5.42 and 7.30 eV, respectively, and a constant Lorentzian broadening with fwhm of 0.5 eV was applied. Bottom row: LR-TDDFT spectra of (d) TiCl₄, (e) SiCl₄ and (f) CrO₂Cl₂ without SOC corrections. Shifts and broadening are the same as for SOC calculations. Experimental data come from ref. 75–77.






Fig. 2 Mean absolute deviation in eV of XAS LR-TDDFT energies with ADMM1 and ADMM2 compared to full HFX, using the (a) pcseg-1/admm-1, (b) pcseg-2/admm-2 and (c) pcseg-3/admm-3 basis sets on the molecules of the CORE65 benchmark set. Wide clear bars are the first excitation energy MAD and thin dark bars the MAD of the differences between the third and first energies ($\omega_3 - \omega_1$).

2.4 Extended systems

Two systems were considered to evaluate the validity of the method in periodic boundary conditions; liquid water and tetrahedral sodium aluminate (NaAlO_2). The former is a relatively sparse and disordered molecular system whereas the second is a much denser well ordered crystal. This difference motivated our choice, as we want to show the applicability of our method in those different regimes.

Fig. 3 shows the LR-TDDFT spectra of a 64 molecules liquid water periodic cell. To take the dynamical nature of the system into account, 50 frames were randomly selected from a 40 000 steps, 20 ps long molecular dynamics simulation at the BLYP+D3⁹¹ level at a temperature of 300 K. A XAS LR-TDDFT calculation was done for each snapshot and the results averaged. Both the MD simulation and the XANES calculations were conducted in periodic boundary conditions. Note that the 1s core electron of all 64 oxygen atoms were excited every time since their individual contributions to the overall spectrum, which are influenced by their local surroundings, may slightly vary. The features of the experimental spectrum are well identifiable in the

simulated one, although not perfectly reproduced. Note that the divergence is likely at least partially due to the approximate structure of DFT water.

The simulated spectra of sodium aluminate can be seen on Fig. 4. The 128 atom model was first geometry optimized at the PBE+D3 level before a XAS LR-TDDFT calculation was performed. Because of symmetry, all aluminium atoms in the system are equivalent, and thus, the excitation from a single Al 1s core state can be considered. For the same reason, only one aluminium was described at the all-electron level whereas pseudopotentials were used for all remaining atoms. The experimental and simulated spectra are in remarkable agreement.

Both calculations were made using the PBEh ($\alpha = 0.45$) hybrid functional, since it is well established that a high percentage of Hartree–Fock exchange is preferable for the description of core excitations.^{5,7,79} Because of the periodic nature of the systems, the truncated Coulomb operator was used for the ground state HFX energy and the exact exchange




Fig. 3 Oxygen K-edge XANES spectrum of liquid water at 300 K. The ADMM-PBEh ($\alpha = 0.45$) functional and all-electron triple-zeta pcseg-2 and admm-2 basis sets were used on all atoms. The calculated LR-TDDFT spectrum was broadened with Lorentzians of fwhm of 0.6 eV and shifted by 5.50 eV to match measurements.⁸⁵




Fig. 4 Aluminium K-edge XANES spectrum of periodic crystalline sodium aluminate. The LR-TDDFT simulation was carried at the ADMM-PBEh ($\alpha = 0.45$) level of theory. The calculated spectrum was rigidly shifted by 23.55 eV and Lorentzian broadened with a constant fwhm of 1.5 eV to match experimental data.⁸⁶ The all-electron triple-zeta pcseg-2 basis set and its admm-2 auxiliary counterpart were used for the single excited Al atom. All other atoms were described with the DZVP-MOLOPT-SR-GTH⁸⁷ and ADMM-FIT3⁷⁰ basis sets and corresponding GTH pseudopotentials.^{88–90}



kernel. As physically sound,³³ the truncation radius was taken to be less than half the simulation cell, namely 6 Å for water and 5 Å for sodium aluminate. All the approximations discussed in the theory part of this paper were used: the sudden approximation, allowing to solve for one core state at a time, the Tamm–Dancoff approximation, which makes the eigenvalue problem Hermitian, the ADMM method, reducing the cost of the ground state hybrid DFT calculation and our XAS specific local RI scheme, which makes the evaluation of the kernel integrals particularly efficient. Fig. 3 and 4 show that the combination of these shortcuts does not undermine the quality of the results.

3 Scaling and performance

Within the Tamm–Dancoff approximation and in periodic boundary conditions, three stages of the XAS LR-TDDFT calculations dominate in terms of cost. First of which is the evaluation of the ADMM four-center electron repulsion integrals. Because of integral screening and the use of short range exchange potentials, this is expected to scale linearly with the size of the system, albeit with a large prefactor. The second stage involves all the exchange correlation kernel integrals of eqn (21), (22) and (24). Thanks to the local RI scheme used there, all these quantities are purely local and expected to require a constant amount of work for each excited atom, for an overall linear scaling with system size. Finally, setting up and diagonalizing the eigenvalue problem of eqn (20) suffers from the worst scaling, as both matrix multiplications (by the projector matrix Q^e) and diagonalization technically scale as $\mathcal{O}(n^3)$. However, the use of atom-centered Gaussian basis sets in large systems tends to bring sparsity into the matrices, reducing the cost of multiplication. Moreover, iterative solvers bring the cost of diagonalization down to $\mathcal{O}(mn^2)$, where m is the number of lowest eigenvalues to solve for. Since the focus is the near edge region of the spectrum, m tends to be small. Note that because the matrix operations take place for each excited atom in the system, the overall scaling is further increased by a power of one.

To actually measure the scaling of the method, calculations were performed on systems of increasing sizes with constant computational power. Because of their different features, both liquid water and crystalline sodium aluminate are considered. The PBEh ($\alpha = 0.45$) functional with truncation radii of 6 and 5 Å was used. All calculations were performed on 12 nodes of a Cray XC50 machine with 12 cores per nodes and a total of 384 GB of RAM. Note that depending on the amount of available memory, all ADMM integrals cannot necessarily be stored and a fraction needs to be recomputed at each SCF cycle. For a fair comparison between the different systems, only the cost of the initial integral evaluation is reported.

Fig. 5 shows the measured costs of XAS LR-TDDFT calculations for water systems of increasing sizes (cubic cells) with the double-zeta pcseg-1 and admm-1 basis sets. Matrix operations are systematically dominating the overall timings and are typically one to two orders of magnitude more expensive than kernel integrals




Fig. 5 Execution times of the most computationally intensive steps of oxygen K-edge XAS LR-TDDFT calculations for liquid water cells of increasing sizes. The ADMM-PBEh ($\alpha = 0.45$) functional with 6 Å truncation radius as well as the pcseg-1 and admm-1 basis sets were used. Excitations from all oxygen atoms take place for each calculation. Least square linear fits of the log–log relations yield polynomial scalings to the power of 0.9, 2.4 and 1.6 for the ADMM ERIs, XAS LR-TDDFT matrix operations and kernel integrals, respectively. All calculations were done on 144 cores of a Cray XC50 machine (12 nodes).

and up to three orders of magnitude costlier than ADMM ERIs. As expected from the discussion above, ADMM integrals scaling was fitted to be linear (power of 0.9). On the other hand, matrix operations scaling was measured as the power of 2.4, which is much more favorable than the worst case scenario of $\mathcal{O}(n^4)$. This is due to matrix sparsity, which is exploited by DBCSR,⁹² CP2K's matrix–matrix multiplication library. Finally, the kernel integrals scaling was measured as the power of 1.6. This is worse than the theoretical linear scaling and the cause of this will be discussed later. However, since the cost is anyway at least ten times less with respect to matrix operations and still with a more favorable scaling, this is not much of a problem.

On Fig. 6, timings for sodium aluminium XAS LR-TDDFT calculations are reported. Note that the starting 128 atoms system has a cubic cell, but larger models are obtained by doubling one or more dimensions of the initial simulation box, leading to non-cubic cells. Similarly, the initial system has only one aluminium atom from which the excitation takes place, but this number doubles every time. The double-zeta pcseg-1 and admm-1 all-electron basis sets were used for excited aluminium atoms and DZVP-MOLOPT-SR-GTH/ADMM-FIT3 bases as well as GTH pseudopotentials were employed for all other atoms. Here, the ADMM ERIs evaluation systematically dominates and its scaling is also measured to be linear (fitted power of 1.1). Matrix operations are almost negligible for smaller systems, but their measured scaling of power 2.4 indicates that they would eventually take over as the main bottleneck. The cost of computing kernel integrals scales linearly as theoretically predicted (fitted power of 0.9) and is systematically two orders of magnitude less than ADMM ERIs evaluation.

The reason why the kernel integrals scaling differs between liquid water and crystalline sodium aluminate is due to both






Fig. 6 Execution times of the computationally dominating parts of aluminium K-edge XAS LR-TDDFT calculations of tetrahedral NaAlO₂ for systems of increasing sizes. Excited aluminium atoms are described with the pceg-1 and admm-1 basis sets and number 1, 2, 4, 8 and 16. All other atoms are described with GTH pseudopotentials and DZVP-MOLOPT-SR-GTH/ADMM-FIT3 basis sets. The functional is ADMM-PBEh ($\alpha = 0.45$) with a 5 Å truncation radius. Least square linear fits of the log–log relations yield polynomial scalings to the power of 1.1, 2.4 and 0.9 for the ADMM ERIs, XAS LR-TDDFT matrix operations and kernel integrals, respectively. All calculations were done on 144 cores of a Cray XC50 machine (12 nodes).

the implementation and the systems differences. While the actual evaluation of the 3-center RI integrals in the AO basis ($pq|\nu$) scales linearly in both cases, it is the dominant operation for NaAlO₂ but not for H₂O. For efficient evaluation of the latter and favorable load balance, multiple optimization steps (integral screening, distribution over the processors, *etc.*) take place and they typically involve loops over the overlap matrix elements (or a fraction of those). Sparsity is such that when the system size doubles, the number of overlap matrix elements increases linearly and so does the cost of optimization, for each excited atom. In the case of water, the combination of high sparsity and small basis sets lead to a special case where the optimization overhead dominates. For sodium aluminate, lower sparsity and larger basis sets (especially the all-electron basis for excited aluminium atoms) lead to negligible optimization costs compared to the actual integral evaluation. A detailed scaling analysis of optimization and integral evaluation costs is available in the supplementary information.

The DBCSR matrix–matrix multiplication library has GPU support implemented. Since matrix operations are so dominant in the water calculations, we investigate whether they can be GPU accelerated. Fig. 7 shows execution times with and without GPU support, for the total calculation and the XAS LR-TDDFT matrix operations only. For the 256 molecules system, the GPU runs are actually slightly slower. This is probably due to the overhead of transferring data to and from the device for relatively few operations. As the systems grow, the GPUs start accelerating the calculations, up to a factor 2.0 in the case of matrix operations for the 1024 molecules system (factor 1.9 for the total execution time).




Fig. 7 Comparison of total execution times (thin dark bars) and XAS LR-TDDFT matrix operations timings (wide clear bars) with and without GPU support for the DBCSR matrix–matrix multiplication library. Functional and basis set choices are the same as for Fig. 5. Calculations were done on 144 cores of a Cray XC50 machine (12 nodes).

4 Conclusions

We have given the details of our LR-TDDFT implementation for core-level spectroscopy in the Sternheimer formalism^{13,14} and using Gaussian type orbital basis sets. Building on existing approaches to XAS LR-TDDFT that decouple core and valence excitations^{6,7} and exploit the sudden approximation,⁸ we further increase the efficiency by introducing new and accurate approximations. Namely, we exploit the localized nature of core orbitals and atom-centered Gaussian basis sets for a tailored resolution of the identity scheme to reduce the effort of calculating the many 2-electron 4-center integrals required. Additionally, we build the LR-TDDFT eigenvalue problem on top of a hybrid DFT ground state calculation that is greatly accelerated by the auxiliary density matrix method⁷⁰ (ADMM) scheme.

We have shown that our local RI scheme does not impact accuracy by reproducing benchmark results at the K-edge.^{78,79} Moreover, we have demonstrated that our treatment of spin–orbit coupling at the L-edge, combining the auxiliary many electron wave function formalism⁶¹ and the quasi-degenerate perturbation approach⁵⁹ can reproduce experimental results. We have also quantified the error introduced by ADMM on excitation energies and concluded that they are negligible on the scale of a full XANES spectrum. Finally, we have shown the applicability of our method for extended systems in periodic boundary conditions by simulating liquid water and tetrahedral sodium aluminate. We measured a favorable sub-cubic scaling of the cost with system size and reported up to factor 2 speedups when using GPUs in calculations where matrix operations dominate.

Although our method is accurate, efficient and able to tackle large systems, it is still necessary to rigidly shift the simulated spectra to match experiments. This makes it more descriptive rather than predictive, although a Δ SCF¹ calculation could be used to calculate the shift *ab-initio*. Such calculations can however be tedious and hard to converge. Therefore, our next efforts will go into the development of a first principles



correction scheme for XAS LR-TDDFT that can be included in the current approach.

Conflicts of interest

There are no conflicts to declare.

Acknowledgements

This work is supported by the MARVEL National Centre for Competency in Research funded by the Swiss National Science Foundation (grant agreement ID 51NF40-182892). Computer time was generously provided by the Swiss National Supercomputing Center (CSCS) under project IDs d110 and s965.

References

- N. A. Besley, A. T. Gilbert and P. M. Gill, *J. Chem. Phys.*, 2009, **130**, 124308.
- D. Jayawardane, C. J. Pickard, L. Brown and M. Payne, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2001, **64**, 115107.
- B. Hetényi, F. De Angelis, P. Giannozzi and R. Car, *J. Chem. Phys.*, 2004, **120**, 8632–8637.
- D. Hait and M. Head-Gordon, *J. Phys. Chem. Lett.*, 2020, **11**, 775–786.
- K. Lopata, B. E. van Kuiken, M. Khalil and N. Govind, *J. Chem. Theory Comput.*, 2012, **8**, 3284–3292.
- M. Stener, G. Fronzoni and M. de Simone, *Chem. Phys. Lett.*, 2003, **373**, 115–123.
- N. A. Besley and A. Noble, *J. Phys. Chem. C*, 2007, **111**, 3333–3340.
- S. DeBeer George, T. Petrenko and F. Neese, *Inorg. Chim. Acta*, 2008, **361**, 965–972.
- S. Coriani and H. Koch, *J. Chem. Phys.*, 2015, **143**, 181103.
- M. L. Vidal, X. Feng, E. Epifanovsky, A. I. Krylov and S. Coriani, *J. Chem. Theory Comput.*, 2019, **15**, 3117–3133.
- P. Norman and A. Dreuw, *Chem. Rev.*, 2018, **118**, 7208–7248.
- T. D. Kühne, M. Iannuzzi, M. Del Ben, V. V. Rybkin, P. Seewald, F. Stein, T. Laino, R. Z. Khaliullin, O. Schütt and F. Schiffmann, *et al.*, *J. Chem. Phys.*, 2020, **152**, 194103.
- J. Hutter, *J. Chem. Phys.*, 2003, **118**, 3928–3934.
- R. Sternheimer, *Phys. Rev.*, 1951, **84**, 244.
- E. Runge and E. K. Gross, *Phys. Rev. Lett.*, 1984, **52**, 997.
- R. van Leeuwen, *Phys. Rev. Lett.*, 1999, **82**, 3863.
- E. Gross and W. Kohn, *Adv. Quantum Chem.*, 1990, **21**, 287–323.
- M. E. Casida, *Recent Advances In Density Functional Methods: (Part I)*, World Scientific, 1995, pp. 155–192.
- A. L. Fetter and J. D. Walecka, *Quantum theory of many-particle systems*, Courier Corporation, 2012.
- S. Hirata and M. Head-Gordon, *Chem. Phys. Lett.*, 1999, **314**, 291–299.
- R. Bauernschmitt and R. Ahlrichs, *Chem. Phys. Lett.*, 1996, **256**, 454–464.
- L. S. Cederbaum, W. Domcke and J. Schirmer, *Phys. Rev. A: At., Mol., Opt. Phys.*, 1980, **22**, 206.
- Y. Zhang, J. D. Biggs, D. Healion, N. Govind and S. Mukamel, *J. Chem. Phys.*, 2012, **137**, 194306.
- W. Liang, S. A. Fischer, M. J. Frisch and X. Li, *J. Chem. Theory Comput.*, 2011, **7**, 3540–3547.
- N. Schmidt, R. Fink and W. Hieringer, *J. Chem. Phys.*, 2010, **133**, 054703.
- J. Rehr, E. Stern, R. Martin and E. Davidson, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1978, **17**, 560.
- J. Foster and S. Boys, *Rev. Mod. Phys.*, 1960, **32**, 300.
- J. Pipek and P. G. Mezey, *J. Chem. Phys.*, 1989, **90**, 4916–4926.
- N. Marzari and D. Vanderbilt, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1997, **56**, 12847.
- O. Vahtras, J. Almlöf and M. Feyereisen, *Chem. Phys. Lett.*, 1993, **213**, 514–518.
- N. A. Besley, *J. Chem. Theory Comput.*, 2016, **12**, 5018–5025.
- Y. Tawada, T. Tsuneda, S. Yanagisawa, T. Yanai and K. Hirao, *J. Chem. Phys.*, 2004, **120**, 8425–8433.
- M. Guidon, J. Hutter and J. VandeVondele, *J. Chem. Theory Comput.*, 2009, **5**, 3010–3021.
- A. Nakata, Y. Imamura and H. Nakai, *J. Chem. Phys.*, 2006, **125**, 064109.
- G. Tu, Z. Rinkevicius, O. Vahtras, H. Ågren, U. Ekström, P. Norman and V. Carravetta, *Phys. Rev. A: At., Mol., Opt. Phys.*, 2007, **76**, 022506.
- Y. Imamura and H. Nakai, *Int. J. Quantum Chem.*, 2007, **107**, 23–29.
- J.-W. Song, M. A. Watson, A. Nakata and K. Hirao, *J. Chem. Phys.*, 2008, **129**, 184113.
- N. A. Besley, M. J. Peach and D. J. Tozer, *Phys. Chem. Chem. Phys.*, 2009, **11**, 10350–10358.
- S. DeBeer George and F. Neese, *Inorg. Chem.*, 2010, **49**, 1849–1853.
- F. A. Asmuruf and N. A. Besley, *Chem. Phys. Lett.*, 2008, **463**, 267–271.
- G. Lippert, J. Hutter and M. Parrinello, *Theor. Chem. Acc.*, 1999, **103**, 124–140.
- G. L. Stoychev, A. A. Auer and F. Neese, *J. Chem. Theory Comput.*, 2017, **13**, 554–562.
- E. F. Valeev, Libint: A library for the evaluation of molecular integrals of many-body operators over Gaussian functions, <http://libint.valeev.net/>, 2020, version 2.7.0-beta.6.
- S. Lehtola, C. Steigemann, M. J. Oliveira and M. A. Marques, *SoftwareX*, 2018, **7**, 1–5.
- B. P. Pritchard, D. Altarawy, B. Didier, T. D. Gibson and T. L. Windus, *J. Chem. Inf. Model.*, 2019, **59**, 4814–4820.
- T. H. Dunning, *J. Chem. Phys.*, 1989, **90**, 1007–1023.
- D. E. Woon and T. H. Dunning, *J. Chem. Phys.*, 1995, **103**, 4572–4585.
- R. Krishnan, J. S. Binkley, R. Seeger and J. A. Pople, *J. Chem. Phys.*, 1980, **72**, 650–654.
- T. Clark, J. Chandrasekhar, G. W. Spitznagel and P. V. R. Schleyer, *J. Comput. Chem.*, 1983, **4**, 294–301.
- A. D. Becke, *Phys. Rev. A: At., Mol., Opt. Phys.*, 1988, **38**, 3098.



- 51 C. Lee, W. Yang and R. G. Parr, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1988, **37**, 785.
- 52 J. P. Perdew, K. Burke and M. Ernzerhof, *Phys. Rev. Lett.*, 1996, **77**, 3865.
- 53 P. J. Stephens, F. J. Devlin, C. F. Chabalowski and M. J. Frisch, *J. Phys. Chem.*, 1994, **98**, 11623–11627.
- 54 C. Adamo and V. Barone, *J. Chem. Phys.*, 1999, **110**, 6158–6170.
- 55 A. D. Becke, *J. Chem. Phys.*, 1993, **98**, 1372–1377.
- 56 J. M. Kasper, P. J. LeStrange, T. F. Stetina and X. Li, *J. Chem. Theory Comput.*, 2018, **14**, 1998–2006.
- 57 G. Fronzoni, M. Stener, P. Decleva, F. Wang, T. Ziegler, E. Van Lenthe and E. Baerends, *Chem. Phys. Lett.*, 2005, **416**, 56–63.
- 58 M. Kadek, L. Konecny, B. Gao, M. Repisky and K. Ruud, *Phys. Chem. Chem. Phys.*, 2015, **17**, 22566–22570.
- 59 M. Roemelt, D. Maganas, S. DeBeer and F. Neese, *J. Chem. Phys.*, 2013, **138**, 204101.
- 60 D. Maganas, S. DeBeer and F. Neese, *J. Phys. Chem. A*, 2018, **122**, 1215–1227.
- 61 F. Franco de Carvalho, B. F. Curchod, T. J. Penfold and I. Tavernelli, *J. Chem. Phys.*, 2014, **140**, 144103.
- 62 E. van Lenthe, J. G. Snijders and E. J. Baerends, *J. Chem. Phys.*, 1996, **105**, 6505–6516.
- 63 C. van Wüllen, *J. Chem. Phys.*, 1998, **109**, 392–399.
- 64 P.-O. Löwdin, *Phys. Rev.*, 1955, **97**, 1474.
- 65 Y. Shao, M. Head-Gordon and A. I. Krylov, *J. Chem. Phys.*, 2003, **118**, 4807–4818.
- 66 F. Wang and T. Ziegler, *J. Chem. Phys.*, 2004, **121**, 12191–12196.
- 67 Z. Li and W. Liu, *J. Chem. Phys.*, 2012, **136**, 024107.
- 68 D. L. Strout and G. E. Scuseria, *J. Chem. Phys.*, 1995, **102**, 8448–8452.
- 69 J. Spencer and A. Alavi, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2008, **77**, 193110.
- 70 M. Guidon, J. Hutter and J. VandeVondele, *J. Chem. Theory Comput.*, 2010, **6**, 2348–2364.
- 71 P. Merlot, R. Izsák, A. Borgoo, T. Kjærgaard, T. Helgaker and S. Reine, *J. Chem. Phys.*, 2014, **141**, 094104.
- 72 E. Rebolini, R. Izsák, S. S. Reine, T. Helgaker and T. B. Pedersen, *J. Chem. Theory Comput.*, 2016, **12**, 3514–3522.
- 73 F. Weigend and R. Ahlrichs, *Phys. Chem. Chem. Phys.*, 2005, **7**, 3297–3305.
- 74 D. Rappoport and F. Furche, *J. Chem. Phys.*, 2010, **133**, 134105.
- 75 A. Wen and A. Hitchcock, *Can. J. Chem.*, 1993, **71**, 1632–1644.
- 76 J. Bozek, K. Tan, G. Bancroft and J. Tse, *Chem. Phys. Lett.*, 1987, **138**, 33–42.
- 77 G. Fronzoni, M. Stener, P. Decleva, M. d. Simone, M. Coreno, P. Franceschi, C. Furlani and K. Prince, *J. Phys. Chem. A*, 2009, **113**, 2914–2925.
- 78 Y. Imamura, T. Otsuka and H. Nakai, *J. Comput. Chem.*, 2007, **28**, 2067–2074.
- 79 P. J. LeStrange, P. D. Nguyen and X. Li, *J. Chem. Theory Comput.*, 2015, **11**, 2994–2999.
- 80 F. Jensen, *J. Chem. Theory Comput.*, 2014, **10**, 1074–1085.
- 81 M. A. Ambroise and F. Jensen, *J. Chem. Theory Comput.*, 2018, **15**, 325–337.
- 82 D. Golze, L. Keller and P. Rinke, *J. Phys. Chem. Lett.*, 2020, **11**, 1840–1847.
- 83 V. Atalla, M. Yoon, F. Caruso, P. Rinke and M. Scheffler, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 2013, **88**, 165122.
- 84 C. Kumar, H. Fliegl, F. Jensen, A. M. Teale, S. Reine and T. Kjærgaard, *Int. J. Quantum Chem.*, 2018, **118**, e25639.
- 85 I. Waluyo, D. Nordlund, U. Bergmann, D. Schlesinger, L. G. Pettersson and A. Nilsson, *J. Chem. Phys.*, 2014, **140**, 244506.
- 86 J. L. Fulton, N. Govind, T. Huthwelker, E. J. Bylaska, A. Vjunov, S. Pin and T. D. Smurthwaite, *J. Phys. Chem. B*, 2015, **119**, 8380–8388.
- 87 J. VandeVondele and J. Hutter, *J. Chem. Phys.*, 2007, **127**, 114105.
- 88 S. Goedecker, M. Teter and J. Hutter, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1996, **54**, 1703.
- 89 C. Hartwigsen, S. Goedecker and J. Hutter, *Phys. Rev. B: Condens. Matter Mater. Phys.*, 1998, **58**, 3641.
- 90 M. Krack, *Theor. Chem. Acc.*, 2005, **114**, 145–152.
- 91 S. Grimme, J. Antony, S. Ehrlich and H. Krieg, *J. Chem. Phys.*, 2010, **132**, 154104.
- 92 U. Borštnik, J. VandeVondele, V. Weber and J. Hutter, *Parallel Computing*, 2014, **40**, 47–58.

